

“Information-Friction” and its implications on minimum energy required for communication

Pulkit Grover, ECE, Carnegie Mellon University

Email: pulkit@cmu.edu

Abstract

Just as there are frictional losses associated with moving masses on a surface, what if there were frictional losses associated with moving information on a substrate? Indeed, many modes of communication suffer from such frictional losses. We propose to model these losses as proportional to “bit-meters,” *i.e.*, the product of mass of information (*i.e.*, the number of bits) and the distance of information transport. We use this “information-friction” model to understand fundamental energy requirements on encoding and decoding in communication circuitry. First, for communication across a binary input AWGN channel, we arrive at fundamental limits on bit-meters (and thus energy consumption) for decoding implementations that have a predetermined input-independent length of messages. For encoding, we relax the fixed-length assumption and derive bounds for flexible-message-length implementations. Using these lower bounds we show that the *total* (transmit + encoding + decoding) energy-per-bit must diverge to infinity as the target error probability is lowered to zero. Further, the closer the communication rate is maintained to the channel capacity (as the target error-probability is lowered to zero), the faster the required decoding energy diverges to infinity.

I. INTRODUCTION




Fig. 1. A Newtonian inspiration for the information-friction model. The units of measuring energy are “bit-meters,” which is the product of number of bits of information, and the Euclidean distance to which that information travels, in the computation.

Just as there are frictional losses associated with moving masses on a surface, there can be frictional losses associated with moving information between gates (see Fig. 1) on a computational substrate. Within the context of communication,

these frictional losses can be a significant part of the energy consumed in computations at the transmitter and the receiver (e.g., encoding and decoding an error-correcting code), which in turn can be a significant fraction of total energy for short-distance communication [1].

What computational models allow us to account for these frictional losses? Communication complexity, introduced by Andrew Yao in [2], accounts for information-movement on a computational substrate by counting the number of bits that need to be moved. However, for many implementations [3] (as discussed in Section IV), energy of computation depends not only on the number of bits, but also on the distance (Euclidean, *i.e.*, L_2 , or “Manhattan” [4], *i.e.*, L_1) to which those bits are moved. Are there models that account for these distances as well?

The VLSI model, introduced by Thompson and others in [5]–[10] (and explored further in [11]–[15]), accounts for these distances by measuring the total wiring infrastructure required to compute a function. The product of the total wiring length and the number of clock-cycles needed, suitably scaled, is used as an approximation for energy consumed in computing. The required wiring infrastructure, as well as energy, are explored through upper and lower bounds (e.g. [6, Ch. 3 and Ch. 4]).

The focus on wires also limits the VLSI model in many ways. First, modern technology is exploring and using alternative interconnects (e.g., optical, carbon nanotubes, or even wireless [16]), and our nervous system uses axons and dendrites, none of which are made of metal wires, and can even evolve (if slowly) as the computation proceeds (e.g. synapses in the brain and wireless interconnects) [17]. Second, modeling computational nodes as ones having small degree of connectivity, as is the case in the VLSI model [6], can be too limiting. Third, even for metal-interconnects, the VLSI model focuses more on the wiring infrastructure needed to move information than on the amount and the distance of information actually moved in the computation. This can overestimate the energy requirements: for instance, not all wires need to be charged and discharged in each clock-cycle, but the model estimates energy consumption based on this assumption¹. Finally, the lengths of messages passed on wires can be different in response to the input of computation, and thus energy-costs can be input dependent. This energy-difference is not accounted for in Thompson’s model.

In Section II-B, we introduce the “information-friction” model of computation and energy consumption (see Fig. 1) that partially addresses these limitations of the VLSI-inspired models. Besides overcoming the limitations addressed above, the model is also appealing because of its conceptual simplicity and fewer assumptions in comparison with the VLSI model. The information-friction model accounts for the cost of computing by counting the “bit-meters”: the product of the number of bits, and the distance to which these bits are moved (summed over all computation links). A similar “bit-meters” metric was used as a measure of “transport capacity” supported by a communication network in the work of Gupta and Kumar [18]. Here,

¹Thompson does acknowledge this shortcoming in his thesis [6].

we are interested in the opposite question: how many bit-meters are *needed to support* a computation?

When is “bit-meters” an appropriate metric for circuit communication energy? The issue is discussed in depth in Section IV, where we argue that for many realistic models of computation (including computation on VLSI circuits), the energy consumption in links in the computational network is well approximated as (or is lowered bounded by) $\mu \times \text{bm}$ where μ is a constant called the *coefficient of information-friction*, and bm are the bit-meters required for the computation. Despite its intuitive appeal and applications, the metric has its shortcomings and limitations, which are also discussed in Section IV.

In Section III, we use the implementation model and an AWGN-based hard-decision channel model to derive the bit-meters cost for decoding an error-correcting code. Intellectually, our work builds on work of El Gamal, Greene, and Pang [19] that uses the VLSI model, to estimate complexity (but not energy) of encoding and decoding an error correcting code. This work also builds on our own work [3] where we derive tradeoffs between wiring area and number of clock-cycles within Thompson’s VLSI-model. In this paper, we show that the required bit-meters for decoding can be no smaller than $\Omega\left(\sqrt{\log \frac{1}{P_e^{blk}}/P_T}\right)$, where P_e^{blk} is the block-error probability, and P_T is the transmit power (for a binary-input AWGN channel where the receiver makes a hard decision on the channel output before decoding, see Section II-A). We show a similar result for encoding under a stronger model of circuit implementation: where the scheduling of messages along the communication links is not predetermined, but can adapt itself to the input of computation. Taking a step further, we also establish that if the communication rate R is maintained close to the channel capacity $C(P_T)$ even as the (block) error-probability P_e^{blk} is driven to zero, the required per-bit energy goes up at least as fast as $\Omega\left(\sqrt{\frac{\log n}{\log \frac{1}{2p_{ch}}}}\right)$. Here, n is the blocklength of the code, and p_{ch} is the cross-over probability of the Binary-Symmetric Channel (BSC) over which the signal is being communicated. As is well known, $n \gtrsim \Omega\left(\frac{\log \frac{1}{P_e^{blk}}}{K(C-R)^2}\right)$ for some constant K (that depends on p_{ch}), and thus diverges to infinity faster as the rate and channel capacity are brought close to each other.

What are the implications of these results on *total* (transmit + computation) energy consumption in communication? Under the information-friction model, optimizing over P_T , we show that the total (transmit + decoding) energy per bit is at least $\Omega\left(\sqrt[3]{\log \frac{1}{P_e^{blk}}}\right)$. This means that for any implementation that experiences information-frictional losses, the total energy per bit must diverge to infinity as the error probability is driven to zero. Further, operating with bounded transmit power (e.g., by operating close to the Shannon limit) appears² to incur larger costs: the total energy per-bit is at least $\Omega\left(\sqrt{\log \frac{1}{P_e^{blk}}}\right)$.

Our results on information-frictional energy for encoding and decoding, and total energy for communication, attempt to begin to fill a void in our understanding of energy required for communication. In a paper that is little-known within the

²In absence of good upper bounds (that are a work in progress), we are left with comparing the lower bounds on energy consumed by the two strategies, which can only offer suggestions on which strategy is more energy-efficient.

information-theory community [20], Landauer argues that one *can* communicate with arbitrarily small energy, paralleling his results on zero-energy reversible computation [21]. In order to do so, however, Landauer observes that one needs to lower friction and noise in the communication medium to effectively zero³, which however requires lowering the speed of computing (asymptotically) to zero to keep the system in thermodynamic equilibrium. From this perspective, information-theoretic works of Golay [30] and Verdú [31] derive capacity per-unit energy for various communication media (*i.e.*, channels) that do have friction and noise, but implicitly assume that computation at the transmitter and receiver is frictionless and noiseless (and hence is free). In this paper, we take a step forward by allowing frictional losses in both communication and computation media and derive lower bounds on energy, whilst still ignoring noise in computation for simplicity.

II. SYSTEM MODEL AND NOTATION

A. Channel model

We consider a point-to-point communication link. An information sequence of k fair coin flips \mathbf{b}_1^k is encoded into 2^{nR} binary-alphabet codewords \mathbf{X}_1^n . The rate of the code is therefore $R = \frac{k}{n}$ bits/channel use, which is assumed to be fixed. The codeword \mathbf{X}_1^n is modulated using BPSK modulation and sent through an Additive White Gaussian Noise (AWGN) channel of bandwidth W , with W channel uses per second. The decoder estimates the input sequence $\hat{\mathbf{b}}_1^k$ by first performing a hard-decision on the received channel symbols before using these (binary) hard-decisions \mathbf{Y}_1^n to decode the input sequence. The overall channel $\mathbf{X}_1^n \rightarrow \mathbf{Y}_1^n$ is therefore a Binary Symmetric Channel (BSC) with raw bit-error probability $p_{ch} := \mathbb{Q}\left(\sqrt{\frac{\zeta P_T}{\sigma_z^2}}\right)$, where $\mathbb{Q}(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt$, ζ is the path-loss associated with the channel, P_T is the transmit power of the BPSK-modulated signal, and σ_z^2 is the variance of the Gaussian noise in the hard-decision estimation. The encoder-channel-decoder system operates at an average block-error probability P_e^{blk} given by $P_e^{blk} = \Pr\left(\hat{\mathbf{b}}_1^k \neq \mathbf{b}_1^k\right)$.

Definition 1 (Channel Model (ζ, σ_z^2)): Channel Model (ζ, σ_z^2) denotes (as described above) a $\text{BSC}(p_{ch})$ channel that is a result of hard-decision at the receiver across an AWGN channel of average transmit power P_T , path loss ζ and noise variance σ_z^2 .

³Of course, from an engineering viewpoint, it makes little sense to think about energy of computing assuming friction and noise are (or can be made) negligible. However, Landauer's main goal was not to provide practically relevant limits to energy of computing (as he himself acknowledges in [20]), but instead to understand and resolve the paradox of Maxwell's demon [22]. This fictional demon is able to lower the thermodynamic entropy of a system seemingly without expending any energy, a violation of the Second Law of Thermodynamics, which would mean (among other "calamitous" conclusions) that perpetual motion machines can exist. A fundamental limit on energy required for communication with arbitrarily small friction and noise would resolve the paradox (because measurement can be viewed as communication of information from the source to the measuring device). Landauer's contention in [20] is that no such limit can exist and thus the paradox cannot be resolved by alluding to energy costs of communication. Instead it is losses in *erasing* information that (according to Landauer) resolve the paradox. We refer the interested reader to [23]–[29] for contemporary work on energy of communication and computing within the context of theoretical physics, and discussions on whether Landauer's principle indeed resolves the paradox.

B. Implementation, computation, and energy models

The computation is performed using a “circuit” on a “substrate.” This section formally defines these terms allowing for decoding analysis in Section III.

Definition 2 (Substrate): A Substrate is a square $Sq(l)$ of side l in \mathbb{R}^2 with vertices at $(0, 0)$, $(0, l)$, $(l, 0)$, and (l, l) .

Definition 3 (SquareLattice(λ)): A SquareLattice(λ) is the collection of points $(s\lambda, t\lambda) \in \mathbb{R}^2$ for all $s, t \in \mathbb{Z}$.

Definition 4 (Grid(λ)): Grid(λ) is the intersection of SquareLattice(λ) with the substrate $Sq(l)$, that is, it is the set of the lattice-points of the square lattice that lie in the substrate.

The parameter λ determines how close computational nodes in the circuit can be brought to each other, and depends on the technology of implementation. For large circuits, $\lambda \ll l$.

Definition 5 (Circuit, computational nodes): The substrate $Sq(l)$ together with a collection $\mathcal{S} \subset \text{Grid}(\lambda)$ of points (called *computational nodes*, or simply *nodes*) inside $Sq(l)$, is called a *Circuit*, and is denoted by $\text{Ckt} = (Sq(l), \mathcal{S})$.

For instance, $Sq(10\lambda)$ along with the set $\mathcal{S} = \{(\lambda, \lambda), (5\lambda, 4\lambda)\}$ constitutes a Circuit.

Nodes can be *input nodes*, *output nodes*, or *helper nodes*. Physically, the nodes help perform the computation by computing functions of received messages. Each node is accompanied with a finite storage memory. Input nodes store the input of computation (one bit each; at the beginning of computation), output nodes store the output (one bit each; at the end of computation), and helper nodes help perform the computation.

Definition 6 (Subcircuit): A subcircuit $\text{SubCkt}_1 = (F_1, \mathcal{S}_1)$ of a circuit $\text{Ckt} = (Sq(l), \mathcal{S})$ is constituted by an open and convex subset F_1 of $Sq(l)$ and by the subset of computational nodes $\mathcal{S}_1 = F_1 \cap \mathcal{S}$.

That is, all the computational nodes within the sub-substrate F_1 must lie in the subcircuit SubCkt_1 .

Definition 7 (Link): A (unidirectional) link connects two nodes in that it allows for noiseless communication between nodes in one direction. The messages are binary-strings. Each message is a function of all the messages (and the possible inputs) received at the transmitting node until the start of the message-transmission.

In a circuit with n nodes, there are $n(n - 1)$ unidirectional links, which can be used more than once during a computation.

Definition 8 (Communication on a circuit): Computational nodes use messages received thus far in computation, and stored memory values, to generate messages that can be communicated to other nodes over links.

We now introduce two models of computation: those with fixed and flexible-length messages. For both, the order of messages passed between computational nodes is pre-determined, but for a flexible-message-length computation, the length of a message can depend on the computation input.

Definition 9 (Fixed-message-length computation (on a circuit)): The computation starts with the arrival of the input of

computation at the input nodes. Each input node stores one bit of the input. The computation then proceeds with communication of messages of predetermined size, *i.e.*, the messages' size does not depend on the input of computation. Each message is a function of the messages that the transmitting computational node has received thus far in the computation (including one bit of the input if the transmitting node is an input node). At the end of the computation, the output is available in the memories of the output nodes.

Definition 10 (Flexible-message-length computation (on a circuit)): The computation is said to be flexible-message-length computation if the number of bits in a message on a link in the computation can depend on the input of computation. Nevertheless, the minimum message-length is assumed to be at least one bit.

A computation may use some or all of the communication links in the circuit. Each link can be used as many times as needed, and at each use, the message can be of any chosen size with the associated costs as described in the following definitions.

Definition 11 (bit-meters cost of a link and of a circuit): The bit-meters cost of a *link* in a computation *Comp* on a circuit *Ckt* is the product of the total number of bits carried by the messages on the link and the Euclidean distance between the nodes at the ends of the link. The bit-meters for the entire circuit *Ckt* is the sum of bit-meters for all the links in *Comp*. Fixing the order of messages (but not necessarily the length), along with making the minimum message-size one bit, makes sure that there's no free-of-cost "silence" [32] that can be used for communicating messages between nodes. Since each message on a link contains at least one bit, and the link is at least λ in length, the message costs at least λ bit-meters.

When a flexible-message-length computation is executed, the bit-meters expended can depend on the input of computation. In such cases, we will often be interested in *average* bit-meters for a link or a computation, where the average is taken over the possible input realizations (with a specified distribution).

Definition 12 (bit-meters for a link within a subcircuit): For a link that connects two nodes within a subcircuit in a computation *Comp*, the bit-meters for that link *within the subcircuit* is the same as the bit-meters for the link in the original circuit. However, if only one of the nodes lies within the subcircuit, then bit-meters for this link within the subcircuit is the product of the number of bits of the message passed along this link and the length of link from the node inside the subcircuit to the boundary of the subcircuit.

Definition 13 (bit-meters for a subcircuit): The bit-meters for a subcircuit $\text{SubCkt}_1 = (F_1, S_1)$ in computation *Comp* is the sum of bit-meters for all the links within the subcircuit (wholly or partially, as defined in Definition 12), and is denoted by $\text{bit-meters}(\text{SubCkt}_1)$.

The definition also holds for bit-meters for the entire circuit.

Definition 14 (Coefficient of information-friction (μ)): The coefficient of information-friction, denoted by μ , characterizes

the energy required for computation in our model. This energy is given by $E = \mu \times \text{bm}$, where bm is the number of bit-meters expended in executing the given computation on a circuit.

Definition 15 (Implementation Model (λ, μ)): Implementation Model (λ, μ) denotes the implementation model as described in this section with λ being the minimum distance between computational nodes, and μ being the coefficient of information-friction.

The same implementation model can be used to execute a fixed or flexible-message-length computation.

III. LOWER BOUNDS ON bit-meters AND INFORMATION-FRICTION ENERGY OF ENCODING AND DECODING

To obtain lower bounds on bit-meters for encoding and decoding, similar to analysis in [3], [19], [33], we need to cut the circuit under consideration into many disjoint subcircuits. The following definitions and lemmas set up the technical background needed for circuit-cutting and ensuing analysis.

Definition 16 (Disjoint subcircuits): Two subcircuits $\text{SubCkt}_1 = (F_1, \mathcal{S}_1)$ and $\text{SubCkt}_2 = (F_2, \mathcal{S}_2)$ of a circuit $\text{Ckt} = (\text{Sq}(l), \mathcal{S})$ are said to be *disjoint subcircuits* if $F_1 \cap F_2 = \phi$, the null set. Similarly, $\{\text{SubCkt}_i\}_{i=1}^{N_{\text{subckt}}}$ are said to be mutually disjoint subcircuits if $F_i \cap F_j = \phi$ for every $i, j \in \{1, 2, \dots, N_{\text{subckt}}\}, i \neq j$.

It follows that any two disjoint subcircuits cannot share computational nodes or communication links that connect two nodes *within* one of the subcircuits. In fact, two disjoint subcircuits do not share bit-meters of computation:

Lemma 1: Let $\{\text{SubCkt}_i\}_{i=1}^{N_{\text{subckt}}}$, where $\text{SubCkt}_i = (F_i, \mathcal{S}_i)$, be a set of mutually disjoint subcircuits of the circuit $\text{Ckt} = (\text{Sq}(l), \mathcal{S})$. Then for any computation Comp ,

$$\text{bit-meters}(\text{Ckt}) \geq \sum_{i=1}^{N_{\text{subckt}}} \text{bit-meters}(\text{SubCkt}_i). \quad (1)$$

Proof: The lemma follows from the observation that in Definition 11, no bit-meters are double-counted in disjoint subcircuits. We note that there are potential situations when $\bigcup_{i=1}^{N_{\text{subckt}}} F_i = \text{Sq}(l)$ for which (1) is not satisfied with equality. This happens when there is a long link in a circuit which has a part that does not lie within either of the subcircuits that contain the two nodes at the ends of the link. ■

The decoder circuit is partitioned into multiple subcircuits via a “Stencil⁴” that can be “moved” over the circuit by changing its origin.

Definition 17 (Stencil): A *Stencil* (a, η, O) in \mathbb{R}^2 , for $\eta < \frac{1}{2}$, is a pattern of equally spaced “inner” squares that are concentric with “outer” squares which form a grid (as shown in Fig. 2). The length of a side of each outer square is a , and the origin O

⁴We use the term “Stencil” in analogy with the classic stencil instrument used to produce letters or designs on an underlying surface. A stencil can be slid on the surface to produce the design at any location on the surface, effectively shifting the origin-point of the design. In this case, a pattern of inner and outer squares is produced on the computational substrate.

lies in the center of an “inner” square. The side of each inner square is of length $s = (1 - 2\eta)a$.

A node in a circuit is said to be *covered* by a Stencil that is overlaid on the circuit substrate if it lies inside an inner-square of the Stencil. For the decoder, the n input nodes store the channel observations, and the k output nodes, also called “bit-nodes,” store the decoded message bits. At the encoder, the k information-bits that are the input of computation are assumed to be stored in bit-nodes. Inside the i -th subcircuit, let k_i^{inside} denote the number of bit-nodes that lie inside the *inner* square, and n_i denote the number of input nodes that lie inside the *outer* square (*i.e.*, anywhere inside the i -th subcircuit).

Definition 18 (Stencil-partition): The outer squares of $\text{Stencil}(a, \eta, O)$ induce a partition (see Fig. 2) of a circuit into subcircuits, each occupying substrate area at most a^2 . If any computational node lies on the boundary of an outer square, then it is arbitrarily included in one of the subcircuits.




Fig. 2. A Stencil overlaid on the Substrate. Also shown are the computational nodes of the Circuit on the Substrate. A zoomed-in version shows the dimensions of the Stencil. As an example, for the square in the zoomed-in version, $k_i^{\text{inside}} = 3$.

The next lemma shows that by moving the Stencil over the substrate, we can find at least one position of the Stencil so that the average number of nodes (over random locations of the Stencil) are covered.

Lemma 2: For any circuit implemented in Implementation Model (λ, μ) , for any $\eta > 0$, there exists an origin O of $\text{Stencil}(a, \eta, O)$ such that the number of bit-nodes covered by the Stencil is lower bounded by

$$\sum_i k_i^{\text{inside}} \geq k(1 - 2\eta)^2. \quad (2)$$

Proof: The proof uses the probabilistic method [34]. Let $O \sim \mathbb{U}\{[0, a), [0, a)\}$, that is, uniformly distributed in the square formed by $(0, 0), (0, a), (a, a), (a, 0)$. Now, the average number of bit-nodes covered by the Stencil (averaged over O) is:

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^k \mathbb{1}_{\{i \text{ covered}\}} \right] &= \sum_{i=1}^k \mathbb{E} [\mathbb{1}_{\{i \text{ covered}\}}] \\ &= \sum_{i=1}^k \Pr(i \text{ covered}) \\ &\stackrel{(a)}{=} \sum_{i=1}^k (1 - 2\eta)^2 \\ &= k(1 - 2\eta)^2 \end{aligned} \tag{3}$$

$$\tag{4}$$

where the key step (a) follows from the observation that for any point, as we move the origin O around uniformly, the probability measure of the set of origins for which the point is covered by the Stencil is the fraction of area covered by the Stencil, which is $(1 - 2\eta)^2$. Thus there exists at least one value of the origin O such that the number of nodes covered is no smaller than the average. ■

Consider the Stencil shown in Fig. 2. The distance between the inner and the outer squares is ηa . B bits are said to be *communicated* from the “transmitting” part of the circuit to the “receiving” part if the values stored in the receiving part are independent of the B bits prior to communication, and the bits can be recovered (in an error-free manner) from the messages *received at the receiving* part during the process of communication. Notice that this definition is looser than the traditional understanding of communication: we do not stipulate that the *stored values* at the receiving part post-communication be able to recover the B bits.

If B bits are communicated from outside an outer square to inside an inner square in a subcircuit, then, intuitively, the bit-meters associated with the subcircuit should be at least $\eta a B$. The following lemma shows this rigorously:

Lemma 3 (bit-meters and average bit-meters in computations): Consider a circuit implemented in Implementation Model (λ, μ) , and any subcircuit SubCkt obtained using the Stencil-partition defined in Definition 18. For communicating B bits of information from outside an outer-square to inside the corresponding inner-square, bit-meters $\geq \eta a B$ for fixed-length messages. Further, even allowing for a flexible-message-length, the *average* bit-meters $\geq \eta a B$. Similarly, for communicating B bits from inside an inner-square to outside the corresponding outer-square, the average bit-meters $\geq \eta a B$.

Proof:

Fixed-length messages: Consider the concentric N_{cut} square-shaped cuts on the sub-circuit-network, starting with the outer square as a cut, with distance λ separating these cuts, as shown in Fig. 3. The cuts end when distance from the inner square is smaller than λ . This remaining distance is denoted by $\alpha\lambda$ for some $\alpha \in [0, 1)$. The inner square is now included as the final N_{cut} -th cut. Except for the inner square, across each cut, each link has to cross at least λ distance.




Fig. 3. Square cuts are made in order to use the cut-set bounding technique. The directed edges show the links along which information flows in the computation. However, the links do not indicate the relative order of information flow during the computation, or the amount of information they carry.

Further, if the number of bits across any cut, which is the summation of bits passed over all links across the cut, is smaller than B , then B bits cannot be delivered to the inner square. Thus across each cut, the total number of bit-meters should be at least $B\lambda$. If N_{cut} is the number of cuts, the total distance for which at least B bits need to travel is at least $(N_{cut} - 2)\lambda + \alpha\lambda$ which is exactly the distance ηa between the inner square and the outer square. Thus, for fixed-message-length computation, $\text{bit-meters}(\text{SubCkt}) \geq \eta a B$.

Flexible-message-length: Flexible-message-length allows for use of variable-length messages on circuit links that can depend on the input of computation. Nevertheless, to code B bits of information using variable-length coding still requires⁵ at least B bits on average [35, Pg. 110]. ■

A. Decoding lower bounds: fixed-length messages

Lemma 4: If at most $\frac{r}{3}$ bits of information is available to obtain an estimate \widehat{M} of a variable M that is distributed uniformly on the set $\mathcal{M} := \{1, 2, \dots, 2^r\}$, r being a positive integer, then $\Pr(\widehat{M} \neq M) \geq \frac{1}{9}$.

Proof: Applying Fano's inequality [35, Pg. 39] to reconstruction of message M , given the available information I of at most $r/3$ bits, the error probability $P_e := \Pr(\widehat{M} \neq M)$ is lower bounded by

$$\begin{aligned}
 P_e \log(|\mathcal{M}| - 1) + h_b(P_e) &\geq H(M|I) \\
 &= H(M) - H(I) + H(I|M) \\
 &\geq H(M) - H(I) \geq r - r/3 = 2r/3,
 \end{aligned} \tag{5}$$

⁵We remind the reader that “silence” can not be used for communication because each message has at least one bit (see Definition 10).

where $h_b(\cdot)$ on the LHS is the binary entropy function. We now consider two cases:

Case 1: $r = 1$: In this case, $|\mathcal{M}| = 2$ and $\log(|\mathcal{M}| - 1) = 0$, and thus from (5),

$$h_b(P_e) \geq \frac{2}{3}. \quad (6)$$

Since $h_b(x) \leq 2\sqrt{x(1-x)} \leq 2\sqrt{x}$ for $x \in (0, 0.5)$ (see, e.g. [36]), $x \geq \frac{(h_b(x))^2}{4}$. From (6), for $r = 1$,

$$P_e \geq \frac{(h_b(P_e))^2}{4} \geq \frac{4}{9 \times 4} = \frac{1}{9}. \quad (7)$$

Case 2: $r \geq 2$: In this case, $|\mathcal{M}| \geq 4$, and thus using a looser form of (5),

$$\begin{aligned} P_e \log(|\mathcal{M}|) + h_b(P_e) &\geq \frac{2r}{3} \\ \Rightarrow P_e \log(|\mathcal{M}|) + 1 &\geq \frac{2r}{3} \\ \Rightarrow P_e &\geq \frac{\frac{2r}{3} - 1}{\log(|\mathcal{M}|)} = \frac{\frac{2r}{3} - 1}{r} = \frac{2}{3} - \frac{1}{r} \\ &\stackrel{(r \geq 2)}{\geq} \frac{2}{3} - \frac{1}{2} = \frac{1}{6} > \frac{1}{9}. \end{aligned}$$

■

We can now connect information-flow in decoding subcircuits to error probability. The following lemma provides a lower bound on the error probability when the number of bit-meters in a subcircuit of the decoder implementation is sufficiently small.

Lemma 5: For any decoder subcircuit SubCkt_i obtained via Stencil-partitioning of Implementation Model (λ, μ) , with $k_i^{\text{inside}} \geq 1$, if $\text{bit-meters}(\text{SubCkt}_i) < \eta a \frac{k_i^{\text{inside}}}{3}$, then $P_e^{\text{blk}} \geq \frac{(2p_{ch})^{n_i}}{9}$.

Proof: From Lemma 3, since the number of bit-meters for the subcircuit is smaller than $\eta a \frac{k_i^{\text{inside}}}{3}$, and the distance between the outer square and the inner square is ηa meters, at most $\frac{k_i^{\text{inside}}}{3}$ bits of information I can be communicated from outside the outer square to inside the inner square.

We first observe that a $\text{BSC}(p_{ch})$ is a stochastically degraded version of a $\text{BEC}(2p_{ch})$. That is, a decoder that receives channel outputs that pass through $\text{BEC}(2p_{ch})$ can simulate a $\text{BSC}(p_{ch})$ channel by randomly assigning the value 0 or 1 to an erased bit, *i.e.* without any increase in bit-meters. Supplying the decoder with outputs of the erasure channel, we examine the event \mathcal{E} when *all* the n_i channel outputs inside the outer square are erased. This event has probability $(2p_{ch})^{n_i}$.

Conditioning on the erasure event \mathcal{E} , let the (block) probability of not recovering all of the bits inside the i -th inner square, denoted by \vec{b}_i^{in} , be $P_{e,i}^{\mathcal{E}}$. From Fano's inequality [35, Pg. 39] applied to reconstructing the message bits $\vec{b}_i^{\text{in}} \in \mathcal{B}_i$, $|\mathcal{B}_i| = 2^{k_i^{\text{inside}}}$, given the communicated information I of entropy at most $k_i^{\text{inside}}/3$ bits,

$$P_{e,i}^{\mathcal{E}} > \frac{1}{9}. \quad (8)$$

Thus, for any $k_i \geq 1$, the (unconditional) error probability for recovering the k_i^{inside} bits correctly is lower bounded by $\frac{(2p_{ch})^{n_i}}{9}$. Since the block-error probability P_e^{blk} for the entire code is larger than the block-error probability in recovering the k_i^{inside} bits in i -th subcircuit, we obtain the lemma. ■

Lemma 6: For the Implementation Model (λ, μ) , for Stencil-partition with outer-squares of side-length a , the maximum number of computational nodes (input, output, or helper) in a subcircuit is upper bounded by

$$N_{nodes} \leq \frac{a^2}{\lambda^2} + 4\frac{a}{\lambda} + 4. \quad (9)$$

Further, if $\frac{a^2}{\lambda^2} \geq 25$,

$$N_{nodes} \leq 2\frac{a^2}{\lambda^2} \quad (10)$$

Proof: The number of nodes in a Stencil cell is approximately $\frac{a^2}{\lambda^2}$. The actual number could however be larger because of boundary effects. On each axis, allowing for one extra node to be included from either side of the square, the number of nodes is (loosely) upper bounded by $(\frac{a}{\lambda} + 2)^2 = \frac{a^2}{\lambda^2} + 4\frac{a}{\lambda} + 4$. Also note that

$$\frac{2a^2}{\lambda^2} - (\frac{a}{\lambda} + 2)^2 = \frac{a^2}{\lambda^2} - 4\frac{a}{\lambda} - 4 = (\frac{a}{\lambda} - 2)^2 - 8,$$

which is positive (in fact, greater than 1) when $\frac{a}{\lambda} \geq 5$, or $\frac{a^2}{\lambda^2} \geq 25$. ■

Theorem 1: For an error correcting code transmitted over a channel with Channel Model (ζ, σ_z^2) and decoded in a decoder circuit DecCkt implemented in Implementation Model (λ, μ) with fixed-message-length implementation that achieves a block-error probability P_e^{blk} , the decoder bit-meters are lower bounded as:

$$\text{bit-meters(DecCkt)} \geq \frac{k}{48\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} \lambda, \quad (11)$$

$$\text{as long as } \log \frac{1}{10P_e^{blk}} > 50 \log \frac{1}{2p_{ch}}. \quad (12)$$

Remark: When condition (12) is violated in the asmyptopia of $P_e^{blk} \rightarrow 0$, i.e., when

$$50 \log \frac{1}{2p_{ch}} \geq \log \frac{1}{10P_e^{blk}}, \quad (13)$$

the transmit power P_T needs to scale at least as fast as $\Omega\left(\log \frac{1}{P_e^{blk}}\right)$. To see this, we use a known bound [37] on the \mathbb{Q} -function,

namely, $\mathbb{Q}(x) \geq \frac{x}{1+x^2} \frac{e^{-x^2/2}}{\sqrt{2\pi}}$:

$$p_{ch} = \mathbb{Q}\left(\sqrt{\frac{\zeta P_T}{\sigma_z^2}}\right) \geq \frac{\sqrt{\frac{\zeta P_T}{\sigma_z^2}}}{1 + \frac{\zeta P_T}{\sigma_z^2}} \frac{e^{-\frac{\zeta P_T}{2\sigma_z^2}}}{\sqrt{2\pi}}. \quad (14)$$

Thus,

$$\ln\left(\frac{1}{p_{ch}}\right) \leq \ln\left(\sqrt{2\pi} \frac{1 + \frac{\zeta P_T}{\sigma_z^2}}{\sqrt{\frac{\zeta P_T}{\sigma_z^2}}}\right) + \frac{\zeta P_T}{2\sigma_z^2} \quad (15)$$

$$\stackrel{(a)}{<} 2\frac{\zeta P_T}{\sigma_z^2}, \text{ if } \frac{\zeta P_T}{\sigma_z^2} \stackrel{(b)}{\geq} 2, \quad (16)$$

where (a) follows from the observation that $\ln(\sqrt{2\pi}) + \ln\left(\frac{1+x}{\sqrt{x}}\right) + \frac{x}{2} < 2x$ for $x \geq 2$ (a fact that can be verified by simply plotting the two sides of the inequality). Further, if condition (b) is not satisfied, then P_T is bounded, and so is p_{ch} , which means that (12) is not violated in the limit $P_e^{blk} \rightarrow 0$. From (a) above and (13), $P_T = \Omega\left(\log \frac{1}{P_e^{blk}}\right)$ under condition (b). This lower bound, which is derived for the case when condition (12) is not satisfied, is larger than our lower bounds on total power when condition (12) is satisfied (Section III-C).

Proof: The outer squares of the Stencil partition the circuit into subcircuits. Let the i -th subcircuit have n_i channel output nodes available within the *outer* square and k_i^{inside} bit-nodes inside the *inner* square. Using Lemma 2, we choose the origin O of the Stencil so that at least $(1 - 2\eta)^2$ fraction of the k bit-nodes are covered by the *inner* squares, i.e.,

$$\sum_i k_i^{\text{inside}} \geq (1 - 2\eta)^2 k. \quad (17)$$

From Lemma 6 choosing Stencil parameter a to be $\frac{1}{\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} \lambda$, under condition (12),

$$\frac{a^2}{\lambda^2} = \frac{1}{2} \frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}} \underset{\text{(under (12))}}{>} \frac{50}{2} = 25.$$

Thus $\frac{a^2}{\lambda^2} \geq 25$. Using Lemma 6, $n_i \leq \frac{2a^2}{\lambda^2} = \frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}$.

From Lemma 5, if bit-meters for any subcircuit are smaller than $\frac{k_i^{\text{inside}}}{3} \eta a$, then the error probability is lower bounded as

$$P_e^{blk} \geq \frac{(2p_{ch})^{n_i}}{9} \geq \frac{1}{9} (2p_{ch})^{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} = \frac{10}{9} P_e^{blk}, \quad (18)$$

which is a contradiction. Thus, for each decoding subcircuit SubCkt_i obtained via the Stencil-partition,

$$\text{bit-meters}(\text{SubCkt}_i) \geq \frac{k_i^{\text{inside}} \eta a}{3}.$$

From Lemma 2, $\sum_i k_i^{\text{inside}} \geq (1 - 2\eta)^2 k$, therefore, using Lemma 1,

$$\begin{aligned} \sum_{i=1}^{N_{\text{subckt}}} \text{bit-meters}(\text{SubCkt}_i) &\geq \frac{(1 - 2\eta)^2 k \eta a}{3} \\ &= \frac{(1 - 2\eta)^2 k \eta}{3\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} \lambda. \end{aligned}$$

Choosing $\eta = \frac{1}{4}$ yields the theorem. ■

B. Encoding lower bounds: fixed and flexible-message-length

Theorem 2: For an error correcting code encoded in a circuit EncCkt that is implemented in Implementation Model (λ, μ) and transmitted over a channel with Channel Model (ζ, σ_z^2) and with block-error probability P_e^{blk} , the encoder average bit-meters (denoted by $\overline{\text{bit-meters}}$) are lower bounded as:

$$\overline{\text{bit-meters}}(\text{EncCkt}) \geq \frac{k}{48\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} \lambda, \quad (19)$$

$$\text{as long as } \log \frac{1}{10P_e^{blk}} > 50 \log \frac{1}{2p_{ch}}, \quad (20)$$

for both fixed and flexible-message-length encoding.

Proof: We directly show the result for flexible-message-length implementations, which subsume fixed-message-length implementations. At the encoder, k input information bits are mapped to n codeword output bits.




Fig. 4. The figure illustrates the definitions of random variables corresponding to bit-nodes and output (codeword) nodes at the encoder. The values $\vec{Y}^{\text{in}}, \vec{Y}^{\text{mid}}, \vec{Y}^{\text{out}}$ are the counterparts of $\vec{X}^{\text{in}}, \vec{X}^{\text{mid}}, \vec{X}^{\text{out}}$ viewed through the channel. It is important to note that they are not based on circuit partitioning at the *decoder*. Indeed, for deriving bounds for the encoder bit-meters, we assume no implementation constraint on the decoder, so it is not even necessary that the decoder is implemented within the Implementation Model of Section II-B.

We again choose the Stencil parameters $a = \frac{1}{\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} \lambda$ and $\eta = \frac{1}{4}$. Focusing on the i -th encoder subcircuit, let n_i denote the number of codeword symbols inside the i -th encoder subcircuit, and let k_i^{inside} denote the input nodes (that store uncoded information) inside the inner square of the subcircuit. Further, for the i -th subcircuit, (dropping subscript i for simplicity) let the information stored in the input nodes inside the inner square be denoted by \vec{b}^{in} , and in those outside the outer square be \vec{b}^{out} . There are more input nodes in the “annulus” between the inner square and the outer square, denote them by \vec{b}^{mid} (see Fig. 4). Similarly, define codeword symbols $\vec{X}^{\text{in}}, \vec{X}^{\text{mid}}, \vec{X}^{\text{out}}$ and the corresponding channel outputs $\vec{Y}^{\text{in}}, \vec{Y}^{\text{mid}}, \vec{Y}^{\text{out}}$ (see Section II-A).

Now, at the decoder, declare the values of \vec{X}^{out} for free. Further, assume that the decoder is not required to recover the values of $\vec{b}^{\text{mid}}, \vec{b}^{\text{out}}$. Thus the job of the decoder is to only recover \vec{b}^{in} (this relaxation on requirements from the decoder will only further reduce the error probability). For recovering \vec{b}^{in} , it has the channel outputs \mathbf{Y}^m , and the freely declared vector \vec{X}^{out} . Using the erasure-channel argument used in decoding lower bounds (Theorem 1), we assume that \mathbf{Z}^m , the outputs of an erasure channel, are available at the decoder as well (which, as far as this theorem is concerned, is free to run the optimal Maximum Likelihood decoding without the constraints of implementation imposed on the encoder). This will only reduce the error probability for the same number of encoding bit-meters. Further, observing that \vec{X}^{out} are available to the decoder, we are interested in minimizing the entropy $H(\vec{b}^{\text{in}} | \vec{X}^{\text{out}}, \mathbf{Y}^m, \mathbf{Z}^m)$, which is the uncertainty at the decoder in the information bits (that are still undeclared, namely the information bits in the i -th encoder subcircuit) given the information available at the

decoder to decode these bits. Examining this uncertainty,

$$H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \mathbf{Y}^m, \mathbf{Z}^m) \stackrel{(a)}{=} H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \mathbf{Z}^m) \quad (21)$$

$$\stackrel{(b)}{=} H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}), \quad (22)$$

where (a) and (b) follow from the Markov chains $\vec{b}^{\text{in}} \rightarrow \{\vec{X}^{\text{out}}, \mathbf{Z}^m\} \rightarrow \mathbf{Y}^m$ and $\vec{b}^{\text{in}} \rightarrow \{\vec{X}^{\text{out}}, \vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} \rightarrow \vec{Z}^{\text{out}}$ respectively.

Similarly,

$$\begin{aligned} & H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \mathbf{Y}^m = \mathbf{y}^m, \mathbf{Z}^m = \mathbf{z}^m) \\ &= H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \mathbf{Z}^m = \mathbf{z}^m) \\ &= H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \vec{Z}^{\text{in}} = \vec{z}^{\text{in}}, \vec{Z}^{\text{mid}} = \vec{z}^{\text{mid}}). \end{aligned} \quad (23)$$

That is, the equality (21) also holds for specific values of the random variables \mathbf{Y}^m and \mathbf{Z}^m .

Our next step, which is key to this proof, is a simple equality. Consider the event that all of the symbols $\{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\}$ are erased, denoted by $\{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} = E$. Then,

$$H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}, \{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} = E) = H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}). \quad (24)$$

This is because the event $\{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} = E$ does not alter the joint distribution of $\vec{b}^{\text{in}}, \vec{X}^{\text{out}}$ *even when encoding is a flexible-message-length computation*. The encoder has no knowledge of this erasure-event⁶, and thus cannot alter the joint distribution in response to the event. Further, under this erasure-event, because $\{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\}$ are completely erased, they provide no help in decoding \vec{b}^{in} .

Thus, if $H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}) \geq \frac{2k_i^{\text{inside}}}{3}$ (as in (5)), then the conditional probability of error in recovering these bits, $\Pr(\vec{b}^{\text{in}} \neq \hat{\vec{b}}^{\text{in}} | \{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} = E)$, is at least $\frac{1}{9}$ (from Lemma 4), and thus the (unconditional) block-error probability is lower bounded by

$$\begin{aligned} P_e^{\text{blk}} &\geq \Pr(\vec{b}^{\text{in}} \neq \hat{\vec{b}}^{\text{in}}) \\ &\geq \Pr(\{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} = E) \Pr(\vec{b}^{\text{in}} \neq \hat{\vec{b}}^{\text{in}} | \{\vec{Z}^{\text{in}}, \vec{Z}^{\text{mid}}\} = E) \\ &\geq \frac{(2p_{\text{ch}})^{n_i}}{9}, \end{aligned} \quad (25)$$

⁶In absence of feedback from the receiver, the encoder only knows the channel statistics, not the realization. While feedback from the receiver to the transmitter is absent here, in presence of noiseless feedback, our bound on encoding $\overline{\text{bit-meters}}$ could be beaten. But the question is more interesting and relevant with realistic models of noisy feedback, where benefits are severely curtailed (see, e.g. [38]). Further, it is also important to note that for flexible-message-length implementations, the key equality (24) holds only when we are investigating circuits at the encoder. At the decoder, the knowledge that all inputs in the subcircuit are erased can be used by a subcircuit to ask for more information from the rest of the decoding circuit. At this point, it is unclear to us if this means that flexible-message-length decoding can beat our bound in Theorem 1.

which leads to a contradiction (following the exact sequence of steps in (18) from proof of Theorem 1).

Thus $H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}) < \frac{2k_i^{\text{inside}}}{3}$ for all i . This means that

$$\begin{aligned} I(\vec{b}^{\text{in}}; \vec{X}^{\text{out}}) &= H(\vec{b}^{\text{in}}) - H(\vec{b}^{\text{in}}|\vec{X}^{\text{out}}) \\ &> k_i^{\text{inside}} - 2k_i^{\text{inside}}/3 = k_i^{\text{inside}}/3. \end{aligned} \quad (26)$$

Thus, at least $k_i^{\text{inside}}/3$ bits of information are communicated from inside the inner square to outside the outer square for each subcircuit i at the encoder. From Lemma 3, the required bit-meters (average or deterministic) for the computation is at least $\eta a \frac{k_i^{\text{inside}}}{3} = \frac{1}{12} k_i^{\text{inside}} a$ (since $\eta = \frac{1}{4}$) for each subcircuit i during encoding, and thus the total average bit-meters for encoding circuitry is at least $\frac{1}{12} k^{\text{inside}} a = \frac{1}{48} k a$, yielding the lemma. ■

We emphasize that while our lower bounds for fixed and flexible-message-length encoding are the same, this does not imply that flexible-message-length cannot reduce the required energy consumption because our bounds could be loose. As we discuss in Section V, this necessitates a comparison with upper bounds, which is a work in progress.

C. Lower bounds on **total** energy consumption

This section uses the bounds on bit-meters derived above to yield bounds on total (transmit and information-friction) energy consumed in communications. Strictly speaking, our bounds are for total energy-per-bit. However these bounds can be translated to total power consumption simply by dividing both transmission and circuit energy by the available time (under the assumption that encoding/decoding can take only as much time as transmission in order to not have buffer-overflows). The results in this section can be viewed as those that account for frictional losses in both the communication channel and the transmitter and receiver circuitry. However, our emphasis is on observing qualitative differences between bounds on total energy and the traditional understanding on transmit energy. Thus we fix the distance (and hence also the path-loss) between the transmitter and the receiver, focusing on the contribution of circuit energy bounds to the total energy.

Corollary 1 (Unavoidable limits on total energy-per-bit): For communication over a channel with Channel Model (ζ, σ_z^2) with the encoder and the decoder implemented in Implementation Model (λ, μ) with fixed-message-length computing, the total energy per bit for communication at error probability P_e^{blk} is lower bounded as:

$$\frac{E_{\text{total}}}{k} \geq \Omega \left(\sqrt[3]{\log \frac{1}{P_e^{\text{blk}}}} \right). \quad (27)$$

Proof: The lower bound considers only the energy at the transmitting end: the transmit and the encoding energy, ignoring the decoding energy. This makes no difference to the order-sense result since the bounds in Theorem 1 and Theorem 2 are the same.

Because the channel is used W times per second, the per-bit transmit energy used is $\frac{nP_T}{W}$. The total (transmit + encoding) energy-per-bit under condition (12) can therefore be lower bounded as (using Theorem 2, and denoting total transmit energy by E_{Tx} , and encoding energy by E_{enc}):

$$\begin{aligned} \frac{E_{total}}{k} &> \frac{E_{Tx} + E_{enc}}{k} \\ &\geq \frac{1}{k} \frac{nP_T}{W} + \frac{1}{k} \frac{\mu k}{48\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}} \\ &= \frac{P_T}{RW} + \frac{\mu}{48\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{\log \frac{1}{2p_{ch}}}}. \end{aligned}$$

In our hard-decision channel model, as P_T increases, the term $\log \frac{1}{2p_{ch}}$ scales proportionally to the received power ζP_T (see, e.g. [3]). Thus

$$\frac{E_{total}}{k} \geq \frac{P_T}{RW} + \frac{\beta}{48\sqrt{2}} \sqrt{\frac{\log \frac{1}{10P_e^{blk}}}{P_T}},$$

for some $\beta > 0$. By simple differentiation, the choice of P_T that minimizes the RHS is $P_T^* = \Theta \left(\sqrt[3]{\log \frac{1}{P_e^{blk}}} \right)$. Substituting,

$$\frac{E_{total}}{k} \geq \Omega \left(\sqrt[3]{\log \frac{1}{P_e^{blk}}} \right). \quad (28)$$

If (12) is not satisfied, then $P_T = \Omega \left(\log \frac{1}{P_e^{blk}} \right)$ (see *Remark* after the statement of Theorem 1), which is larger than the behavior in (28). ■

Remark: While these bounds hold for any fixed communication distance in the limit of $P_e^{blk} \rightarrow 0$, it is important to note that for practically interesting values of P_e^{blk} (typically between 10^{-3} and 10^{-20}), empirical evidence [39]–[41] suggests that relative to transmit power, circuit power is relevant only at short distances (less than a few kilometers). At longer distances, the energy consumed in circuits at high P_e^{blk} can be neglected in total power optimization because the transmit power is dominant. However, there can be situations where decoding power is still important because the receiver can be more energy constrained than the transmitter (e.g. in the downlink of a cellular system).

D. What happens as the code-rate approaches the channel capacity?

In practical situations, transmit power can be constrained by regulating authorities (e.g. the FCC) or the limit of the power amplifier at the transmitter circuitry. In such situations, it is not possible to increase transmit power to reduce the required encoding and decoding power. While our past work has shown that energy can be expended in other components (e.g. the equalizer or the beamformer) to effectively increase the SNR at the decoder [42], thereby providing analogous tradeoffs between transmit and circuit power as above, there likely are saturation-effects to such approaches as well (e.g. the thermal noise limit or interference due to ambient transmissions that are unaccounted for).

What happens when the code rate is maintained near channel capacity (or, by keeping transmit power near Shannon limit for a fixed rate, the channel capacity is maintained near the code rate) even in the asymptotic limit of $P_e^{blk} \rightarrow 0$? Is the energy-cost higher than the case when we relax the constraint of operating close to capacity? Our earlier work shows this is the case [1] for energy consumed in *computational nodes* in the VLSI model (but does not show it for wiring energy, or the information-frictional energy for movement of information). Is this the case for information-frictional energy as well? The theorem below proves that this is indeed the case, and in fact, the information-frictional energy consumption is significantly higher (in order sense) than the energy consumed in computational nodes. The key observation used in the derivation of the following result is that small enough bit-meters in computation can lead to multiple sub-circuits having local decoding errors due to independent channel events. Because error in any one subcircuit leads to a block-error, and the error-events used to lower bound the error-probability of different subcircuits are independent, a stronger lower bound can be derived that captures a stronger dependence on n .

Theorem 3: For an error correcting code transmitted over a channel with Channel Model (ζ, σ_z^2) and decoded in a decoder circuit DecCkt implemented in Implementation Model (λ, μ) with fixed-message-length implementation and block-error probability P_e^{blk} , the decoder bit-meters are lower bounded as:

$$\text{bit-meters} \geq \frac{k}{192} \sqrt{\frac{\log n}{\log \frac{1}{2^{p_{ch}}}}} \lambda, \quad (29)$$

as long as

$$\log n > 100 \log \frac{1}{2^{p_{ch}}} \quad (30)$$

Proof: See Appendix A. ■

Remark: The theorem shows that (under condition (30)) as $n \rightarrow \infty$, the required bit-meters per-bit, *i.e.* $\frac{\text{bit-meters}}{k}$ diverge to infinity as $\sqrt{\log n}$ for fixed transmit power. It is well known (e.g. [43, Exercise 5.23] [44]) that close to capacity, as P_e^{blk} is made small for a fixed rate, $n \gtrsim \Omega\left(\frac{\log \frac{1}{P_e^{blk}}}{K(C-R)^2}\right)$ for some constant K (that depends on p_{ch}). That is, the “speed” of increase of block length (and hence also of bit-meters per-bit) as $P_e^{blk} \rightarrow 0$ blows up as the code-rate approaches capacity.

Further, note that condition (30) is satisfied in the asymptotic limit $P_e^{blk} \rightarrow 0$ for fixed-rate communication problems where communication is close to capacity. This is because in such situations, the transmit power needs to be maintained close to the Shannon limit (a constant at fixed rate), and thus $\log \frac{1}{p_{ch}}$ is bounded even as $P_e^{blk} \rightarrow 0$.

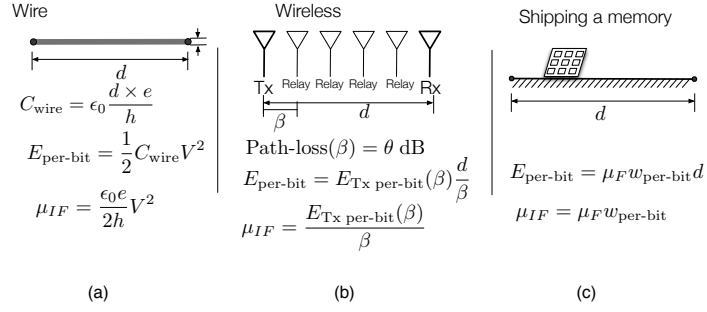


Fig. 5. Example ways of communication where the Information-Friction Model is a good approximation for energy consumption. The coefficient of information-friction, μ_{IF} , is also calculated. For clarity, we use μ_{IF} for coefficient of information-friction here to distinguish it from μ_F , the Newtonian coefficient of friction. In (a), C_{wire} is the capacitance of an on-chip wire of length d , width e , and distance h from the computational substrate. ϵ_0 is the permittivity of air. In (b), β is the minimum possible distance between two repeaters, and $E_{\text{Tx per-bit}}(\beta)$ is the required energy per-bit to communicate to distance β . In (c), $w_{\text{per-bit}}$ is the weight of a memory cell that stores a bit.

IV. JUSTIFICATION FOR, AND THE LIMITATIONS OF, THE INFORMATION-FRICTION MODEL

A. Practical examples where information-friction model applies

The following modes of communicating (via binary symbols⁷) in a computational system motivate our definition of information-friction:

Metal wires: A metal wire can be modeled as a capacitance between the wire and the substrate (often referred to as the Elmore lumped model [4]) that needs to be charged and discharged through the wire resistance (see Fig. 5 (a)). The amount of energy expended by a single charge/discharge operation (e.g., to communicate one bit of information) is $\frac{1}{2}C_{\text{wire}}V^2$ [4], which increases linearly with the wire-capacitance (here V is the voltage across the wire). Further, the capacitance $C_{\text{wire}} = \epsilon \frac{d \times e}{h}$ itself increases linearly with the wire-length d (here e is the wire-width, and h is the distance of the wire from the substrate). Thus each charge/discharge requires energy that scales linearly with the wire-length. The resulting “coefficient of friction” is shown in Fig. 5 (a).

Wireless links: While wireless communication from a single transmitting node to a single receiving node suffers from worse-than-linear losses (path-loss is often $\frac{1}{d^\psi}$ for some $\psi > 2$), with repeaters placed at uniform distances, the energy requirement can be brought down to approximately linear with distance. However, this is only possible when the signal-to-noise ratio is large, which is when the errors are so improbable that their propagation along the relays can be neglected. When these errors cannot be neglected, or when the repeaters are difficult to place, the information-friction model only provides a loose bound on the total energy.

Transporting matter: One way of communicating is via writing the message on a memory, and sending the memory from

⁷The binary-symbol-assumption is made for simplicity. The results can easily be extended for any fixed “constellation size.”

the sender to the receiver [45]. Not only is this form of communication widely used today (e.g. use of USB sticks and CDs to transport information), it has been envisaged as a method of communication in a billiard-balls computer by the physicists [46], and has also been thought of as an efficient way of communicating across interstellar space [45]. Any friction (e.g. friction between the medium of transport, such as the billiard ball, and the surface) that the transported object faces gets translated into information-friction via the weight of one bit of storage unit. It appears that pneumatic computers (that store and communicate between logic elements using fluid movements, e.g. [47]) would encounter similar frictional losses. For fluid traveling through smooth pipes, there is still a loss in pressure which is linear per-unit length (see [48]). For communicating computational messages reliably, this pressure loss will again necessitate use of repeaters, thereby leading to information-frictional losses just as those for wireless links.

B. Limitations of the information-friction model

At extremely low speeds of computation, it may be possible to reduce the coefficient of information-friction, consistent with results in thermodynamics of computation [21], [49], [50], by communicating using timing of the signal (such as in pulse-position modulation, or through silence [32]). In such situations, with a single change in the message on a circuit-link, a large number of bits can be communicated (depending on how slow the computation can be). However, such techniques are hard to implement because they require sophisticated synchronization between circuit components in order to exploit communication via timing. Often this synchronization is performed by explicitly sending a clock-signal [4], and the communication of clock-signal itself can consume significant amount of energy. Thus it is unclear if communication using timing is a practical way to reduce the coefficient of information-friction significantly.

While in most situations, information-friction bounds are valid (if loose) lower bounds on energy-consumption, we note that there could be situations where these bounds are beaten. One such situation is when a computation uses wireless broadcast for transmission on computation links. It is plausible, for instance, that when multicasting to multiple nodes simultaneously, the required energy can increase slower than linearly with the cumulative distance of communication. There is literature that uses broadcast as a way to reduce communication requirements in the sense of traditional (Andrew Yao's) communication complexity of distributed sorting [51]. A deeper exploration is needed to understand if energy requirements can also be lowered for such computations via broadcast to beat the information-friction limits.

Finally, we note that information-frictional energy is not always the dominant sink of energy in computational systems. While asymptotically, our theoretical results here and empirical observations in [52], [53] strongly suggest that information-friction is the dominant sink, in practical systems, energy consumed in computational nodes or memory-access could be significant, and could even dominate in non-asymptotic scenarios. Improved modeling of energy consumed in nodes and memory-access

could enhance the understanding in such scenarios.

V. DISCUSSIONS AND CONCLUSIONS

The information-friction model proposed here can be viewed as a broadening and a simplification of the VLSI model introduced by Thompson and others. The model enjoys several advantages over the VLSI model. In particular, it can capture energy requirements in wired as well as wireless computational systems, and has a closer connection to energy consumption (as noted in the introduction). Within information-theoretic literature, our metric of bit-meters for computational costs has been used earlier as a metric for transport capacity of wireless networks [18]. Within physics, it has a potential connection with thermodynamics. Most of the classical analysis focuses on energy of single operations (e.g. [20], [54]), and even this analysis becomes difficult when the computation needs to be performed in non-infinite time⁸, in part because friction can no longer be ignored⁹. Recent works [27]–[29] have shown promise towards addressing finite-time single-operation computing, but even once this is understood, it will still remain to extend the analysis to multi-operation computation. While our techniques here are guided strongly by current implementations, they could complement the single-operation-based analysis in statistical physics, offering suggestions regarding what form the fundamental limits should look like.

Nevertheless, we do believe that an even broader approach is needed to understand how physically-fundamental our limits on energy are. The approach proposed here is not in the spirit of Landauer’s, where the goal is to relax all constraints (timing of computation, frictional energy, medium of implementation, etc.) in obtaining fundamental limits. Instead, this approach is closer to Shannon’s engineering approach: just as Shannon modeled the communication channel and derived fundamental limits that hold for all possible communication strategies *for the chosen channel model*, here we model the communication channel *and the implementation*, and derive limits that hold for all possible communication strategies and implementation architectures and algorithms *for the chosen implementation model*. The key assumptions lie in modeling of implementation, and a good first step towards deeper understanding can be to relax or modify these assumptions¹⁰.

Are these limits useful in guiding code-design? Our complementary work with experimentalists [39], [41] that provides upper bounds on energy has shown that the code-choice needs to adapt to distance of communication: at shorter distances, simpler coding techniques (that require smaller wire-length per-bit) are more total-energy-efficient than capacity-approaching codes. As distances of communication increase, approaching capacity becomes increasingly efficient. In this paper, for reasons

⁸Finite-time analyses need to tackle non-equilibrium thermodynamics, which has proven to be quite hard (e.g. [29]).

⁹Friction can be ignored in infinite-time analysis because changes can be made at speeds approaching zero, keeping the system in equilibrium at all times, lowering frictional losses to as low as desired.

¹⁰As Norbert Wiener noted on choice of assumptions, “What most experimenters take for granted before they begin their experiments is infinitely more interesting than any results to which their experiments lead.”

of clarity, we have fixed the communication distance (see Section II-A). Even for purely intellectual reasons, it is important to explore these upper bounds further and obtain an order-sense asymptotic understanding (along the lines of [52]) on how tight the lower bounds are, and if the suggestions we draw via comparison of lower bounds (e.g. using bounded transmit power as $P_e^{blk} \rightarrow 0$ fundamentally requires larger total power) in this paper actually hold.

One also needs to understand the implications in multi-user situations, especially in interference-limited situations where the advantage of increasing transmit power indefinitely can be limited by saturation of SINR, as explored in [1]. Intuitively in such situations [1], as the density of transmitting devices increases, it becomes increasingly important to save transmit power (that can cause interference) even at the cost of increased encoding and decoding energy. It might be the case that energy-efficient radios need to be “cognitive” in detecting nearby transmitter and receiver density, and choosing the optimal energy-efficient strategy in response.

Finally, an important question remains to be understood in the total energy of point-to-point communication: how much can feedback help? Perfect (noiseless, infinite-precision) feedback can help in reducing complexity significantly [55]. However, perfect feedback is impossible to obtain in practice, and more reliable feedback also requires an increased energy cost (just as more reliable forward transmission does). One will therefore need to examine the issue in presence of noisy feedback, of which the understanding is far from mature, especially from a fundamental-limits perspective (e.g. [38], [56]–[59]). More broadly, we also need to allow noise in the computation process itself (some of our recent work, e.g. [60], [61], focuses on this issue), a line of work started by von Neumann [62] that still lacks a strong connection with energy consumption.

ACKNOWLEDGMENTS

We acknowledge the generous support of NSF grants NSF-ECCS-1343324, NSF CAREER (NSF-CCF-1350314), and a startup grant from Carnegie Mellon University. We also thank Manoj Gopalakrishnan and Sanjoy Mitter for helpful discussions, the reviewers for their detailed reading and helpful suggestions, and Majid Mahzoon and Yaoqing Yang for carefully reading and commenting on the final version.

APPENDIX A

INCREASE IN DECODING ENERGY ON APPROACHING CAPACITY

This Appendix provides the proof of Theorem 3.

Proof: Choose the Stencil parameter $a = \frac{1}{2} \sqrt{\frac{\log n}{\log \frac{1}{2p_{ch}}}} \lambda$ for some $\xi < 1$. Then, under condition (30) (which guarantees that $\frac{a^2}{\lambda^2} = \frac{\log n}{4 \log \frac{1}{2p_{ch}}} > \frac{100 \log \frac{1}{2p_{ch}}}{4 \log \frac{1}{2p_{ch}}} = 25$, satisfying the condition of Lemma 6), by Lemma 6, $n_i \leq \frac{2a^2}{\lambda^2} = \frac{\log n}{2 \log \frac{1}{2p_{ch}}}$.

The rest of the proof uses ideas from the work of Blake and Kschischang [63] to bound block-error probability under independent subcircuit error events, and is via contradiction. Choose $\eta = \frac{1}{4}$, and suppose bit-meters $< \frac{k}{192} \sqrt{\frac{\log n}{\log \frac{1}{2p_{ch}}}} \lambda = \frac{1}{24} k \eta a \leq \frac{1}{24} k^{\text{inside}} a$ (for appropriately chosen Stencil origin). Under this assumption, we first claim (and prove via contradiction) that for at least $\frac{k^{\text{inside}}}{2}$ bit-nodes, the subcircuits that they lie in have bit-meters $_i \leq \frac{k_i^{\text{inside}}}{12} a$. Suppose our claim is not correct. Then for at least $\frac{k^{\text{inside}}}{2}$ bits, the subcircuits they lie in have bit-meters $_i > \frac{k_i^{\text{inside}}}{12} a$, which would mean that the total number of bit-meters is larger than $\frac{k^{\text{inside}}}{24} a$, leading to a contradiction. Thus at least $\frac{k^{\text{inside}}}{2}$ bit-nodes lie in subcircuits with bit-meters $_i \leq \frac{k_i^{\text{inside}}}{12} a$. With $\eta = \frac{1}{4}$, this means that at most $\frac{k_i^{\text{inside}}}{3}$ bits of information is available to decode these k_i^{inside} bits in the event of erasure of all the channel outputs inside the outer square of the i -th subcircuit, leading to a lower bound of $\frac{1}{9}$ on error probability conditioned on this erasure event.

Now notice that at the decoder, these erasure events are independent across different circuits. Further, the information inside *every* subcircuit needs to be recovered in order to recover the entire block. This yields the following stronger lower bound on the block-error probability.

$$P_e^{\text{blk}} \geq 1 - \prod_{i: \text{bit-meters}_i \leq \frac{k_i^{\text{inside}}}{12} a} \left(1 - \frac{(2p_{ch})^{n_i}}{9} \right). \quad (31)$$

where the set $Err := \{i : \text{bit-meters}_i \leq \frac{k_i^{\text{inside}}}{12} a\}$ is the set of subscript-indices such that each such subcircuit has error probability in recovering its information bits lower bounded by $\frac{(2p_{ch})^{n_i}}{9}$. Because at least $\frac{k^{\text{inside}}}{2}$ number of bits lie in subcircuits with bit-meters $_i \leq \frac{k_i^{\text{inside}}}{12} a$, and from Lemma 6, $k_i^{\text{inside}} \leq \frac{2a^2}{\lambda^2} = \frac{\log n}{2 \log \frac{1}{2p_{ch}}}$ for any subcircuit, it has to be the case that

$$|Err| \geq \frac{\frac{k^{\text{inside}}}{2}}{\frac{\log n}{2 \log \frac{1}{2p_{ch}}}} \stackrel{(a)}{\geq} \frac{\frac{k}{4}}{\log \frac{1}{2p_{ch}}} = \frac{nR \log \frac{1}{2p_{ch}}}{4 \log n}, \quad (32)$$

where (a) uses the fact that $k_i^{\text{inside}} \geq (1 - 2\eta)^2 k = \frac{k}{4}$ (since $\eta = \frac{1}{4}$). Thus,

$$P_e^{\text{blk}} \geq 1 - \left(1 - \frac{(2p_{ch})^{\bar{n}}}{9} \right)^{\frac{nR \log \frac{1}{2p_{ch}}}{4 \log n}}, \quad (33)$$

where $\bar{n} := \frac{\log n}{2 \log \frac{1}{2p_{ch}}}$ is also an upper bound on n_i for each i . Examining the second term in the RHS of (33) by taking its

log,

$$\begin{aligned}
& \log \left(1 - \frac{(2p_{ch})^{\bar{n}}}{9} \right)^{\frac{nR \log \frac{1}{2p_{ch}}}{4 \log n}} \\
&= \frac{nR \log \frac{1}{2p_{ch}}}{4 \log n} \log \left(1 - \frac{(2p_{ch})^{\bar{n}}}{9} \right) \\
&= \frac{nR \log \frac{1}{2p_{ch}}}{4 \log n} \log \left(1 - \frac{(2p_{ch})^{\frac{\log n}{2 \log \frac{1}{2p_{ch}}}}}{9} \right) \\
&= \frac{nR \log \frac{1}{2p_{ch}}}{4 \log n} \log \left(1 - \frac{1}{9n^{\frac{1}{2}}} \right) \\
&\approx \frac{nR \log \frac{1}{2p_{ch}}}{4 \log n} \left(-\frac{1}{9\sqrt{n}} \right) \\
&\xrightarrow{n \rightarrow \infty} -\infty.
\end{aligned}$$

Thus, the second term in the RHS of (33) goes to 0, and $P_e^{blk} \rightarrow 1$ as $n \rightarrow \infty$, leading to a contradiction. ■

REFERENCES

- [1] P. Grover, K. Woyach, and A. Sahai, “Towards a communication-theoretic understanding of system-level power consumption,” *IEEE Journal on Selected Areas in Communication*, vol. 29, no. 8, pp. 1744–1755, Sep. 2011.
- [2] A. C.-C. Yao, “Some complexity questions related to distributive computing (preliminary report),” in *ACM symposium on Theory of computing (STOC)*, New York, NY, USA, 1979, pp. 209–213.
- [3] P. Grover, A. Goldsmith, and A. Sahai, “Fundamental limits on the power consumption of encoding and decoding,” in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Jul. 2012, pp. 2716–2720.
- [4] J. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital integrated circuits*. Prentice Hall Englewood Cliffs, New Jersey, 2002.
- [5] C. D. Thompson, “Area-time complexity for VLSI,” in *Proceedings of the 11th annual ACM symposium on Theory of computing (STOC)*. New York, NY, USA: ACM, 1979, pp. 81–88.
- [6] —, “A complexity theory for VLSI,” Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, 1980.
- [7] R. P. Brent and H. Kung, “The area-time complexity of binary multiplication,” *Journal of the ACM*, vol. 28, no. 3, pp. 521–534, 1981.
- [8] B. Chazelle and L. Monier, “Towards more realistic models of computation for VLSI,” in *Proceedings of Caltech Conference on VLSI*, Pasadena, CA, 1981, pp. 441–453.
- [9] C. E. Leiserson, “Area-efficient VLSI computation,” Ph.D. dissertation, Carnegie Mellon University, 1981.
- [10] C. Mead and L. Conway, *Introduction to VLSI systems*. Addison-Wesley Publishing Co., 1980.
- [11] B. P. Sinha and P. K. Srimani, “A new parallel multiplication algorithm and its VLSI implementation,” in *Proceedings of the 1988 ACM 16th annual conference on computer science*. New York, NY, USA: ACM, 1988, pp. 366–372.
- [12] M. R. Kramer and J. van Leeuwen, “The VLSI complexity of boolean functions,” in *Proceedings of the Symposium “Rekursive Kombinatorik” on Logic and Machines: Decision Problems and Complexity*. London, UK: Springer-Verlag, 1984, pp. 397–407.
- [13] S. N. Bhatt, G. Bilardi, and G. Pucci, “Area-time tradeoffs for universal VLSI circuits,” *Theoretical Computer Science*, vol. 408, no. 2-3, pp. 143 – 150, 2008, excursions in Algorithmics: A Collection of Papers in Honor of Franco P. Preparata.
- [14] R. Cole and A. Siegel, “Optimal VLSI circuits for sorting,” *J. ACM*, vol. 35, no. 4, pp. 777–809, 1988.
- [15] C. D. Thompson, “The VLSI complexity of sorting,” *IEEE Transactions on Computers*, vol. 100, no. 12, pp. 1171–1184, 1983.

- [16] M. F. Chang, V. P. Roychowdhury, L. Zhang, H. Shin, and Y. Qian, "RF/wireless interconnect for inter-and intra-chip communications," *Proceedings of the IEEE*, vol. 89, no. 4, pp. 456–466, 2001.
- [17] P. Dayan and L. F. Abbott, *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Taylor & Francis, 2001.
- [18] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.
- [19] A. El Gamal, J. Greene, and K. Pang, "VLSI complexity of coding," in *The MIT Conf. on Adv. Research in VLSI*, Cambridge, MA, Jan. 1984.
- [20] R. Landauer, "Minimal energy requirements in communication," *Science*, vol. 272, no. 5270, pp. 1914–1918, 1996.
- [21] —, "Computation: A fundamental physical view," *Physica Scripta*, vol. 35, pp. 88–95, 1987.
- [22] H. S. Leff and A. F. Rex, *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing*. Institute of Physics, 2003.
- [23] J. Ladyman and K. Robertson, "Landauer defended: Reply to Norton," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2013.
- [24] J. D. Norton, "Author's Reply to Landauer Defended," *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 2013.
- [25] —, "All Shook Up: Fluctuations, Maxwells Demon and the Thermodynamics of Computation," *prepared for Entropy*, 2013.
- [26] M. Hemmo and O. Shenker, "Entropy and Computation: The Landauer-Bennett Thesis Reexamined," *Entropy*, vol. 15, no. 8, pp. 3297–3311, 2013.
- [27] M. Gopalkrishnan, "The Hot Bit I: The Szilard-Landauer Correspondence," *arXiv preprint arXiv:1311.3533*, Nov. 2013.
- [28] —, "The Hot Bit II," *Arxiv preprint in preparation*, Dec. 2013.
- [29] E. Aurell, K. Gawedzki, C. Mejía-Monasterio, R. Mohayaee, and P. Muratore-Ginanneschi, "Refined second law of thermodynamics for fast random processes," *Journal of statistical physics*, vol. 147, no. 3, pp. 487–505, 2012.
- [30] M. J. E. Golay, "Note on the theoretical efficiency of information reception with PPM," in *Proc. IRE*, Sep. 1949, p. 1031.
- [31] S. Verdú, "On channel capacity per unit cost," *IEEE Trans. Inf. Theory*, vol. 36, no. 9, pp. 1019–1030, Sep. 1990.
- [32] A. K. Dhulipala, C. Fragouli, and A. Orlitsky, "Silence-based communication," *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 350–366, 2010.
- [33] P. Grover, "Fundamental limits on the power consumption for lossless signal recovery," in *IEEE Inf. Theory Workshop (ITW)*, Sep. 2012.
- [34] N. Alon and J. H. Spencer, *The Probabilistic Method*, 2nd ed. New York, NY: Wiley-Interscience, 2000.
- [35] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York: Wiley, 2006.
- [36] T. Courtade, "Properties of the binary entropy function," Oct. 2012, blog entry. [Online]. Available: <https://blogs.princeton.edu/blogit/2012/10/26/properties-of-the-binary-entropy-function/>
- [37] R. F. Tate, "On a double inequality of the normal distribution," *The Annals of Mathematical Statistics*, pp. 132–134, 1953.
- [38] C. Li, "Capacity-achieving codes for noisy feedback channels: A necessary condition," *IEEE Communications Letters*, vol. 18, no. 6, pp. 1007–1010, June 2014.
- [39] K. Ganesan, P. Grover, and J. M. Rabaey, "The power cost of overdesigning codes," in *IEEE Workshop on Signal Processing Systems (SiPS)*, Oct. 2011.
- [40] S. L. Howard, C. Schlegel, and K. Iniewski, "Error control coding in low-power wireless sensor networks: when is ECC energy-efficient?" *EURASIP Journal on Wireless Communications and Networking*, pp. 1–14, 2006.
- [41] K. Ganesan, Y. Wen, P. Grover, A. Goldsmith, and J. Rabaey, "Choosing "green" codes by simulation-based modeling of implementations," in *IEEE Globecom*, Dec 2012.
- [42] P. Grover, A. Sahai, and J.-H. Park, "Simple models for power optimization across transmission, equalization and decoding," in *Proceedings of the 19th European Signal Processing Conference (EUSIPCO)*, Barcelona, Spain, Aug. 2011.
- [43] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY: John Wiley, 1971.
- [44] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Dispersion of Gaussian channels," in *IEEE International Symposium on Information Theory*, Seoul, Korea, 2009.

- [45] C. Rose and G. Wright, "Inscribed matter as an energy-efficient means of communication with an extraterrestrial civilization," *Nature*, vol. 431, no. 7004, pp. 47–49, 2004.
- [46] E. Fredkin and T. Toffoli, "Conservative logic," *Int. J. Theoretical Physics*, vol. 21, pp. 219–253, 1982.
- [47] M. Rhee and M. A. Burns, "Microfluidic pneumatic logic circuits and digital pneumatic microprocessors for integrated microfluidic systems," *Lab on a Chip*, vol. 9, no. 21, pp. 3131–3143, 2009.
- [48] B. McKeon, C. Swanson, M. Zagarola, R. Donnelly, and A. J. SMITS, "Friction factors for smooth pipe flow," *Journal of Fluid Mechanics*, vol. 511, pp. 41–44, 2004.
- [49] C. H. Bennett, "Logical reversibility of computation," *IBM journal of Research and Development*, vol. 17, no. 6, pp. 525–532, 1973.
- [50] —, "Notes on Landauer's principle, reversible computation, and Maxwell's Demon," *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, vol. 34, no. 3, pp. 501–510, 2003.
- [51] R. Dechter and L. Kleinrock, "Broadcast communications and distributed algorithms," *IEEE Transactions on Computers*, vol. 100, no. 3, pp. 210–219, 1986.
- [52] K. Ganesan, P. Grover, and A. Goldsmith, "How far are LDPC codes from fundamental limits on total power consumption?" in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2012.
- [53] K. Ganesan, Y. Wen, P. Grover, A. Goldsmith, and J. Rabaey, "Minimizing transmit and decoding power for LDPC codes: Complexity bounds and circuit simulations," *In Preparation for submission to IEEE Trans. Wireless*, 2014.
- [54] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM journal of research and development*, vol. 5, no. 3, pp. 183–191, 1961.
- [55] J. P. M. Schalkwijk and T. Kailath, "A coding scheme for additive noise channels with feedback – I: No bandwidth constraint," *IEEE Trans. Inf. Theory*, vol. 12, no. 2, pp. 172–182, Apr. 1966.
- [56] Y. Kim, A. Lapidoth, and T. Weissman, "Error exponents for the gaussian channel with active noisy feedback," *Information Theory, IEEE Transactions on*, vol. 57, no. 3, pp. 1223–1236, 2011.
- [57] Y. Xiang and Y.-H. Kim, "On the awgn channel with noisy feedback and peak energy constraint," in *IEEE International Symposium on Information Theory (ISIT)*, 2010, pp. 256–259.
- [58] M. V. Burnashev and H. Yamamoto, "On the zero-rate error exponent for a bsc with noisy feedback," *Problems of Information Transmission*, vol. 44, no. 3, pp. 198–213, 2008.
- [59] —, "On the reliability function for a noisy feedback gaussian channel: Zero rate," *Problems of Information Transmission*, vol. 48, no. 3, pp. 199–216, 2012.
- [60] P. Grover, "Is "shannon-capacity of noisy computing" zero?" in *IEEE International Symposium on Information Theory (ISIT)*, 2014, pp. 2854–2858.
- [61] Y. Yang, P. Grover, and S. Kar, "Can a noisy encoder be used to communicate reliably?" in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, to appear, Monticello, IL, Oct. 2014.
- [62] J. von Neumann, "Probabilistic logics and the synthesis of reliable organisms from unreliable components," in *Automata Studies*. Princeton University Press, 1956, pp. 329–378.
- [63] C. Blake and F. R. Kschischang, "Energy of decoding algorithms," in *13th Canadian Workshop on Information Theory (CWIT)*, 2013, pp. 1–5.